

ARE WE THERE YET? TIME TO DETECTION OF NANOHERTZ GRAVITATIONAL WAVES BASED ON PULSAR-TIMING ARRAY LIMITS

S. R. Taylor^{1,2*}, M. Vallisneri^{1,2}, J. A. Ellis^{1,2,†}, C. M. F. Mingarelli^{2, 3, 1}, T. J. W. Lazio^{1,2}, & R. van Haasteren^{1,2,†}

Draft version March 9, 2016

ABSTRACT

Decade-long timing observations of arrays of millisecond pulsars have placed highly constraining upper limits on the amplitude of the nanohertz gravitational-wave stochastic signal from the mergers of supermassive black-hole binaries ($\sim 10^{-15}$ strain at $f = 1/\text{yr}$). These limits suggest that binary merger rates have been overestimated, or that environmental influences from nuclear gas or stars accelerate orbital decay, reducing the gravitational-wave signal at the lowest, most sensitive frequencies. This prompts the question whether nanohertz gravitational waves are likely to be detected in the near future. In this letter, we answer this question quantitatively using simple statistical estimates, deriving the range of true signal amplitudes that are compatible with current upper limits, and computing expected detection probabilities as a function of observation time. We conclude that small arrays consisting of the pulsars with the least timing noise, which yield the tightest upper limits, have discouraging prospects of making a detection in the next two decades. By contrast, we find large arrays are crucial to detection because the quadrupolar spatial correlations induced by gravitational waves can be well sampled by many pulsar pairs. Indeed, timing programs which monitor a large and expanding set of pulsars have an $\sim 80\%$ probability of detecting gravitational waves within the next ten years, under assumptions on merger rates and environmental influences ranging from optimistic to conservative. Even in the extreme case where 90% of binaries stall before merger and environmental coupling effects diminish low-frequency gravitational-wave power, detection is delayed by at most a few years.

Subject headings: methods: data analysis — gravitational waves — pulsars: individual — pulsars: general

1. INTRODUCTION

For the last ten years, three international collaborations have been collecting precise timing observations of the most stable millisecond pulsars, with the aim of identifying the imprint of gravitational waves (GWs) in the nanohertz frequency band (Hobbs 2013; Kramer & Champion 2013; McLaughlin 2013). The best-motivated GW source for such pulsar-timing arrays (PTAs) is the stochastic background signal from the cosmological population of gravitationally bound supermassive black-hole binaries (SMBHBs: Rajagopal & Romani 1995; Jaffe & Backer 2003; Wyithe & Loeb 2003) that are expected to form at the centers of galaxies after these merge (Sesana et al. 2004, 2008). If the binary inspirals are driven purely by GW radiation reaction at orbital separations corresponding to the PTA frequency band, the resulting timing-residual signal has a power-law spectrum with a characteristic exponent,

$$S(f) = \frac{h_c(f)^2}{12\pi^2 f^3} = \frac{A_{\text{GW}}^2}{12\pi^2} \left(\frac{f}{\text{yr}^{-1}} \right)^{-13/3} \text{yr}^3, \quad (1)$$

where $h_c(f)$ is the characteristic strain spectrum of the GW background. Theoretical expectations for the GW amplitude A_{GW} center around a few 10^{-15} (McWilliams et al. 2014; Sesana 2013; Ravi et al. 2014). Recent observational upper limits from the three international pulsar-timing consortia (Lentati et al. 2015; Arzoumanian et al. 2015a; Shannon et al. 2015) imply $A_{\text{GW}} \lesssim 10^{-15}$ at 95% confidence. These limits appear in tension with theoretical expectations, so much so that it seems plausible that environmental effects in galactic centers either stall the formation of gravitationally bound SMBHBs, or accelerate their inspiral; in both cases, the effective A_{GW} is reduced at the frequencies where PTAs are most sensitive (Arzoumanian et al. 2015a; Shannon et al. 2015).

Indeed, the PTA detection of the stochastic GW background (GWB) from SMBHBs, considered imminent only a few years ago (Siemens et al. 2013; McWilliams et al. 2014), now seems to be receding toward the future. Is this really the case? In this letter, we answer this question quantitatively. Namely, given the upper limit $A_{\text{ul}} = 10^{-15}$ obtained by the Parkes PTA (PPTA; Shannon et al. 2015), we ask when we can expect to make a positive detection using different PTAs: the PPTA (Hobbs 2013), limited to the four low-timing-noise pulsars used for the upper limit; the North American Nanohertz Observatory for Gravitational-waves (NANOGrav; McLaughlin 2013); the European PTA (EPTA; Kramer & Champion 2013); the International PTA (IPTA; see, e.g., Hobbs et al. 2010); and the pulsar-timing project that will be supported by SKA1 (Janssen et al. 2015).

* email: Stephen.R.Taylor@jpl.nasa.gov

¹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

² TAPIR Group, MC 350-17, California Institute of Technology, Pasadena, California 91125, USA

³ Max Planck Institute for Radio Astronomy, Auf dem Hügel 69, D-53121 Bonn, Germany

[†] Einstein Fellow

We adopt the frequentist formalism developed by Rosado et al. (2015), and we characterize the PTAs simply by listing, for each pulsar, the duration of the observation and the levels of measurement and timing noise. While GW searches and upper limits with PTAs have recently been given a Bayesian treatment (van Haasteren et al. 2009; van Haasteren & Levin 2013; Lentati et al. 2013), the frequentist approach based on optimal statistics is both convenient and appropriate for the purpose of this paper, because it dispenses with the simulation of actual datasets, and because the quantification of detection probability is intrinsically a frequentist statement (see, e.g., Vallisneri 2012). Furthermore, experience shows that frequentist upper limits and detection prospects are rather close to Bayesian results (see, e.g., Arzoumanian et al. 2015a).

We proceed as follows: for a specified PTA and true GW background amplitude (A_{true}), we compute the detection probability (DP) as a function of observation time, using the detection statistic described below. Detection is a probabilistic endeavor since the actual realization of measurement and timing noise may obscure or expose the underlying GW signal. Indeed, we do not currently have access to A_{true} , but only to the observed upper limit A_{ul} . Setting upper limits is also a probabilistic endeavor, because different realizations of noise lead to different A_{ul} given the same A_{true} , so we use the upper-limit optimal statistic (also described below) to compute $p(A_{\text{ul}}|A_{\text{true}})$. By introducing an astrophysically motivated prior $p(A_{\text{true}})$, we can then use Bayes' theorem to obtain $p(A_{\text{true}}|A_{\text{ul}}) \propto p(A_{\text{ul}}|A_{\text{true}})p(A_{\text{true}})$. Finally, we obtain the expected detection probability $\text{DP}(A_{\text{ul}})$ as $\int \text{DP}(A_{\text{true}})p(A_{\text{true}}|A_{\text{ul}})dA_{\text{true}}$.⁶

We conclude that, while small arrays of a few low-timing-noise pulsars (represented here by the Shannon et al. 2015 configuration) provide the most constraining upper limits on the nanohertz GWB, they are suboptimal for detection in comparison to larger arrays such as NANOGrav, EPTA, full-PPTA, and the IPTA, which are expanded regularly with newly discovered pulsars. Indeed, detection and upper limits have rather different demands, and the PPTA itself may employ a ~ 20 -pulsar array for detection in the near future (Reardon et al. 2016). Encouragingly, we find that larger PTAs have $\sim 80\%$ probability of making a detection in the next ten years, even allowing for a reduced GWB signal due to significant binary stalling or environmental influences.

2. STATISTICS

The details of the frequency-domain statistical framework used here can be found in Rosado et al. (2015). In the frequentist context, a statistic X is a single number that summarizes the (noisy) data with respect to the presence of a (GW) signal, so that, on average, a higher A_{true} yields a higher X , with fluctuations due to the range of possible noise realizations. Given a set of data, there is one observed \hat{X} , and setting an upper limit $A_{\text{ul}}^{95\%}$, amounts to stating that if A_{true} were equal to $A_{\text{ul}}^{95\%}$, for 95% of noise realizations the observed X would have been higher than

⁶ By assuming a prior, we are introducing a Bayesian element in a frequentist scheme, but this is necessary because we wish to make a statement about a range of astrophysical possibilities, which are all compatible to varying degrees with the observed upper limit.

\hat{X} .

Detection schemes based on statistics, on the other hand, proceed as follows: one considers separately the case where a signal is present in the data at a certain level, and the case where it is not, and computes the probability distribution of the statistic (over the ensemble of noise realizations) in both cases. The zero-signal distribution is used to set a detection threshold as a function of a false-alarm probability; the signal distribution is used to compute the probability that for a certain A_{true} the statistic will exceed the threshold—that is, the detection probability (or efficiency).

In the following we use cross-correlation statistics for both upper-limit and detection considerations, making use of the correlated influence of a GWB signal on pulsars which are widely separated on the sky.

Upper limits.— For the discussion of upper limits in this work we will use the so-called A -statistic of Rosado et al. (2015) modified so that the standard deviation is measured in units of squared GWB amplitude and that the value of the statistic is equal to the injected GWB amplitude on average, as is done in Chamberlin et al. (2015). Assuming that the distribution of this statistic is Gaussian (a reasonable approximation at low signal-to-noise ratios (Chamberlin et al. 2015)) we obtain

$$p(A^2|A_{\text{true}}) = \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{(A^2 - A_{\text{true}}^2)^2}{2\sigma_0^2}\right) \quad (2)$$

with

$$\sigma_0 = A_{\text{true}}^{-2} \left(2 \sum_k \sum_{ij} \Gamma_{ij}^2 \frac{S_{h0}^2}{P_i P_j} \right)^{1/2}, \quad (3)$$

where \sum_k denotes a sum over frequencies f_k , \sum_{ij} denotes a sum over pulsar pairs, Γ_{ij} is the overlap reduction function (i.e. this is the Hellings & Downs (1983) curve for an isotropic GWB) for pulsars i and j , $S_{h0} \equiv S_{h0}(f_k)$ is the modeled cross-power spectral density, and $P_i \equiv P_i(f_k)$ is the intrinsic noise power spectral density. From this Gaussian probability distribution, it can be shown that

$$A_{\text{ul}}^2 = \hat{A}^2 + \sqrt{2}\sigma_0 \text{erfc}^{-1}(2(1-C)), \quad (4)$$

where \hat{A} is the measured value of the GWB amplitude and C is our upper limit confidence (e.g., 95%). Thus, the distribution $p(A_{\text{ul}}^2|A_{\text{true}})$ is a Gaussian distribution with standard deviation σ_0 and mean $A_{\text{true}}^2 + \sqrt{2}\sigma_0 \text{erfc}^{-1}(2(1-C))$. Note, to compute $p(A_{\text{ul}}|A_{\text{true}})$ we perform a coordinate transformation to obtain $p(A_{\text{ul}}|A_{\text{true}}) = 2A_{\text{ul}}p(A_{\text{ul}}^2|A_{\text{true}})$.

Detection.— In the previous section we used the A -statistic for upper limits in order to compare our results with upper limits computed using the statistic presented in Chamberlin et al. (2015). For assessing our detection probability, however, we will make use of the B -statistic of Rosado et al. (2015). In this case, we wish to determine the detection probability

$$\text{DP} = \frac{1}{2} \text{erfc} \left[\frac{\sqrt{2}\sigma_0 \text{erfc}^{-1}(2\alpha_0) - \mu_1}{\sqrt{2}\sigma_1} \right], \quad (5)$$

where α_0 is the false alarm probability (FAP); μ_1 , σ_1 , and




Figure 1. The normalized probability distribution for the SMBHB GW amplitude A_{true} , given upper limits from the PPTA (Shannon et al. 2015) and a large PTA which regularly adds pulsars (specifically NANOGrav (Arzoumanian et al. 2015a)). The darker blue curves assume an amplitude prior based on Sesana (2013); the faded red curves modify it by assuming 90% binary stalling, reducing the amplitude $\sqrt{10}$ -fold. The dashed curves reflect turnover spectra due to binary stellar hardening. The dash-dotted curves correspond to the astrophysically-motivated priors on the amplitude for no stalling (blue) and 90% stalling (red).

σ_0 (distinct from previous) are the mean in the presence of a signal, the standard deviation in the presence (subscript 1) and the standard deviation in the absence (subscript 0) of a signal (see Eqs. A16–A18 of Rosado et al. 2015 for more details). Both μ_1 and σ_1 grow with increasing signal amplitude. Here we choose a FAP of 0.13% to match the 3-sigma detection threshold of Siemens et al. (2013).⁷

Implementation.— In the work of Rosado et al. (2015), it is argued that timing model subtraction (namely subtraction of the quadratic term of the timing model; due to pulsar spin down) can be emulated by not including the lowest frequency in the sum over k frequencies. However, we have found that these analytic results agree much better with simulations including timing model subtraction if all frequencies are included in the sum. Furthermore, to emulate data sets with different time spans we only include frequencies in the sum that are greater than $1/T_{\text{min}}$, where T_{min} is the shortest time span for the given pulsar pair.

In this work we always include the GWB power spectral density in the intrinsic pulsar noise P_i . This is meant to mimic a real optimal statistic analysis where the intrinsic noise would be based on single-pulsar-analyses in which there is no way to distinguish the intrinsic noise from GWB power. This will make our analysis conservative in the sense that we are overestimating the intrinsic noise.

3. RESULTS & DISCUSSION

⁷ This choice is clear if one notices that the 3-sigma range is a two-sided confidence limit, thus 0.27% of the probability density function is outside of this range. To convert to a standard FAP, we need the one-sided limit which is simply $0.27/2 \sim 0.13$.

In the following, we characterize detection prospects in terms of small and large PTAs, and specifically consider the following five configurations:

- (1) PPTA4, consisting of the four pulsars described in Shannon et al. (2015) with associated measurement (white) and timing (red) noise (Shannon et al. 2015; Reardon et al. 2016).
- (2) NANOGrav+, consisting of the 37 pulsars described in Table 3 of Arzoumanian et al. (2015b) with their associated noise properties, to be expanded in future observations by adding four new pulsars each year with 250-ns TOA measurement noise, in line with current expectations.
- (3) EPTA+, consisting of the 42 pulsars described in Table 3 of Caballero et al. (2015) with their associated noise properties, and regularly adding pulsars as in (2). Some of these additional pulsars are monitored with the Large European Array for Pulsars (LEAP; Kramer & Champion 2013; Bassa et al. 2016), which synthesizes a 194-m steerable dish from the five EPTA telescopes.
- (4) A conservative IPTA+, consisting of the 49 pulsars of the first IPTA data release, described in Verbiest et al. (2016) with measurement and timing noise⁸ properties as in Lentati et al. (2016), and expanding by six new pulsars each year, again with 250-ns TOA measurement noise.
- (5) A theorist’s PTA (TPTA), consisting of the toy configuration of Rosado et al. (2015) which may be supported by an advanced radio telescope such as SKA1: 50 pulsars with 100-ns TOA measurement noise and no intrinsic timing noise.

These configurations are summarized in Table 1.9

We consider four combined estimates for the amplitude and spectral shape of the stochastic GWB from SMBHBs, ranging from more optimistic (detection wise) to more conservative. For the amplitude, we adopt the observationally motivated log-normal distribution of Sesana (2013) with mean $\log_{10} A_{\text{true}} = -15$ and standard deviation of 0.22; however, we also consider the case where 90% of binaries require longer than ~ 9 –13.7 Gyr before reaching the PTA band (i.e., the binaries stall, Burke-Spolaor & Simon 2015; McWilliams et al. 2014), which corresponds to a mean $\log_{10} A_{\text{true}} = -15.5$, with the same standard deviation. For each of these amplitude priors, we examine the purely GW-driven power-law spectrum of Eq. (1), as well as the case where the GW strain has a turnover at $f = 1/(11 \text{ yr})$ caused by SMBHB interactions with stars in galactic nuclei, which accelerate binary inspiral and remove low-frequency power with respect to the pure power law (see, e.g., Sampson et al. 2015 and Arzoumanian et al. 2015a). The $1/(11 \text{ yr})$ turnover fre-

⁸ For IPTA pulsars timing noise is taken to consist only of red spin noise, and not of any other system- or band-specific red noise. The latter components may be isolated with multi-system and multi-frequency observations, while the former is completely conflated with a GWB signal in the absence of cross-correlation measurements.

⁹ We rescale PPTA, NANOGrav, EPTA, and IPTA measurement noises (but not the timing noise) by a common factor for each PTA, so that the corresponding A -statistic upper limits match the results of Shannon et al. (2015), Arzoumanian et al. (2015a), Lentati et al. (2015), and Verbiest et al. (2016); this rescaling has little impact on detection probability in the future, since white measurement noise becomes subdominant to red timing noise at low frequencies.

Table 1
PTAs considered, as characterized in our analysis.

PTA	configuration	number of pulsars
PPTA4	PPTA (Shannon et al. 2015)	4
NANOGrav+	NANOGrav (Arzoumanian et al. 2015b)	37 + 4 new/year with 250-ns TOA measurement precision
EPTA+	EPTA (Caballero et al. 2015)	42 + 4 new/year with 250-ns TOA measurement precision
IPTA+	IPTA Data Release 1 (Verbiest et al. 2016) (Lentati et al. 2016)	49 + 6 new/year with 250-ns TOA measurement precision
Theorist's PTA 100-ns TOA measurement precision		50

quency was chosen to lie just beyond the low-frequency reach of the current PPTA dataset, since the Shannon et al. (2015) analysis still found $\sim 9\%$ consistency with the purely GW-driven model of Sesana (2013) for $f \geq 1/(11 \text{ yr})$. Similarly, the recent NANOGrav (Arzoumanian et al. 2015a) analysis finds 20% consistency using a 9-year dataset (Arzoumanian et al. 2015b). Furthermore, our choice of a turnover at $1/(11 \text{ yr})$ corresponds to reasonable assumptions about the density of stellar populations interacting with SMBHBs in galactic nuclei (see Figs. 10 and 11 of Arzoumanian et al. 2015a). At very low frequencies, the frequency scaling of the timing-residual spectrum becomes $S(f) \propto 1/f$ due to the environmental influences.

In the top panel of Fig. 1 we show the probability distribution $p(A_{\text{true}}|A_{\text{ul}})$ obtained by setting the 95% A-statistic upper limit to the Shannon et al. (2015) value, 1×10^{-15} , and by assuming no-stalling (solid blue line) and 90%-stalling (faded solid red line) amplitude priors, with pure power-law true spectra. The dashed lines correspond to turnover spectra.¹⁰ For comparison, in the bottom panel of Fig. 1 we show also $p(A_{\text{true}}|A_{\text{ul}})$ as computed for NANOGrav's optimal-statistic upper limit 1.3×10^{-15} (Arzoumanian et al. 2015a). The underlying astrophysically-motivated priors for no-stalling and 90%-stalling are shown in each panel as blue and red dash-dot lines, respectively.

Next, we compute the B-statistic GW detection probabilities, Eq. (5), for each PTA, as a function of A_{true} and of the observation time T beyond current datasets. The resulting DPs are shown in Figs. 2: the left and right columns correspond to the pure power-law and turnover spectra, respectively.¹¹

As mentioned above, for NANOGrav+/EPTA+ and IPTA+ we add new pulsars at rates of four and six per year, respectively. This annual expansion has a positive impact on the DP, since it provides more independent pulsar pairs at differing angular separations, which help discriminate GWs with quadrupolar spatial correlations (the Hellings & Downs 1983 curve, or the more general anisotropic signatures discussed by Mingarelli et al. 2013

and Gair et al. 2014) from non-spatially-correlated red noise processes (Siemens et al. 2013). The unique GW correlation signature is the smoking gun for detection, and we must coordinate future efforts to maximize its measurement.

We obtain the expected detection probabilities $\langle \text{DP} \rangle_{A_{\text{true}}}$ as a function of time by integrating $\text{DP}(A_{\text{true}}, T)$ against the $p(A_{\text{true}}|A_{\text{ul}})$ curves of Fig. 1 (top panel). The resulting $\langle \text{DP} \rangle_{A_{\text{true}}}$ are shown in Figs. 3. As in Fig. 1, the darker blue and lighter red curves correspond to no-stalling and 90%-stalling amplitude priors. For the PPTA4 array, expected detection probabilities remain below 10% throughout the next 20 years of observation, even in the most optimistic GW-signal scenario. This is unsurprising—a PTA consisting of a few exquisitely timed pulsars may provide very tight upper limits, but it will not usually yield convincing detection statistics, which require an array of many pulsars to map out the expected spatial correlation signature.

By contrast, large pulsar arrays such as NANOGrav+, EPTA+, and IPTA+ provide high detection probabilities even with strong binary environmental couplings, since they allow the quadrupolar spatial correlation signature of a GWB to be mapped by many different pulsar pairings. We expect these results to be qualitatively the same for a full PPTA that regularly adds pulsars to the array. In NANOGrav+, EPTA+, and IPTA+, detection probabilities begin to grow rapidly after only five years of observation beyond current datasets. Binary stalling (the red curves) stunts this growth by three years at most. While the low-frequency turnover reduces detection probability, its effect is mitigated by the large number of pulsars and the annual catalog expansion.¹²

Finally, we see that the large number of well-timed pulsars in the TPTA builds highly convincing detection probabilities after only a few years of operation. By the time the pulsars have been observed long enough that the influence of the turnover at $f = 1/(11 \text{ yr})$ may be noticeable, the DP is already close to unity. The same is true were we to add timing noise at currently known lev-

¹⁰ When analyzing these latter cases, we use the “wrong” spectrum – a pure power law – for the P_i in Eq. (3), for consistency with existing analyses that assume GW-driven inspirals in deriving limits. As discussed in Rosado et al. (2015), doing so has negligible effects on the distribution of the A-statistic.

¹¹ In computing the B-statistic, we do assume that the analysis accounts correctly for the true spectral shape. Again, this choice has minimal effects.

¹² Shannon et al. (2015) advocate performing observations with higher cadence to improve sensitivity at $f \geq 0.2 \text{ yr}^{-1}$, where the GWB would be less affected by environmental couplings. In fact, higher-cadence observations would improve sensitivity across all frequencies, but not if they come at the cost of reducing the number of monitored pulsars because of limited observing-time allocations. Furthermore, GW searches are actually sensitive to the spectrum of timing residuals rather than of strain itself. The two are related by $S_h(f) \propto h_c(f)^2/f^3$, so even a turnover spectrum will leave a steep red-noise signature in the pulsar timing residuals.

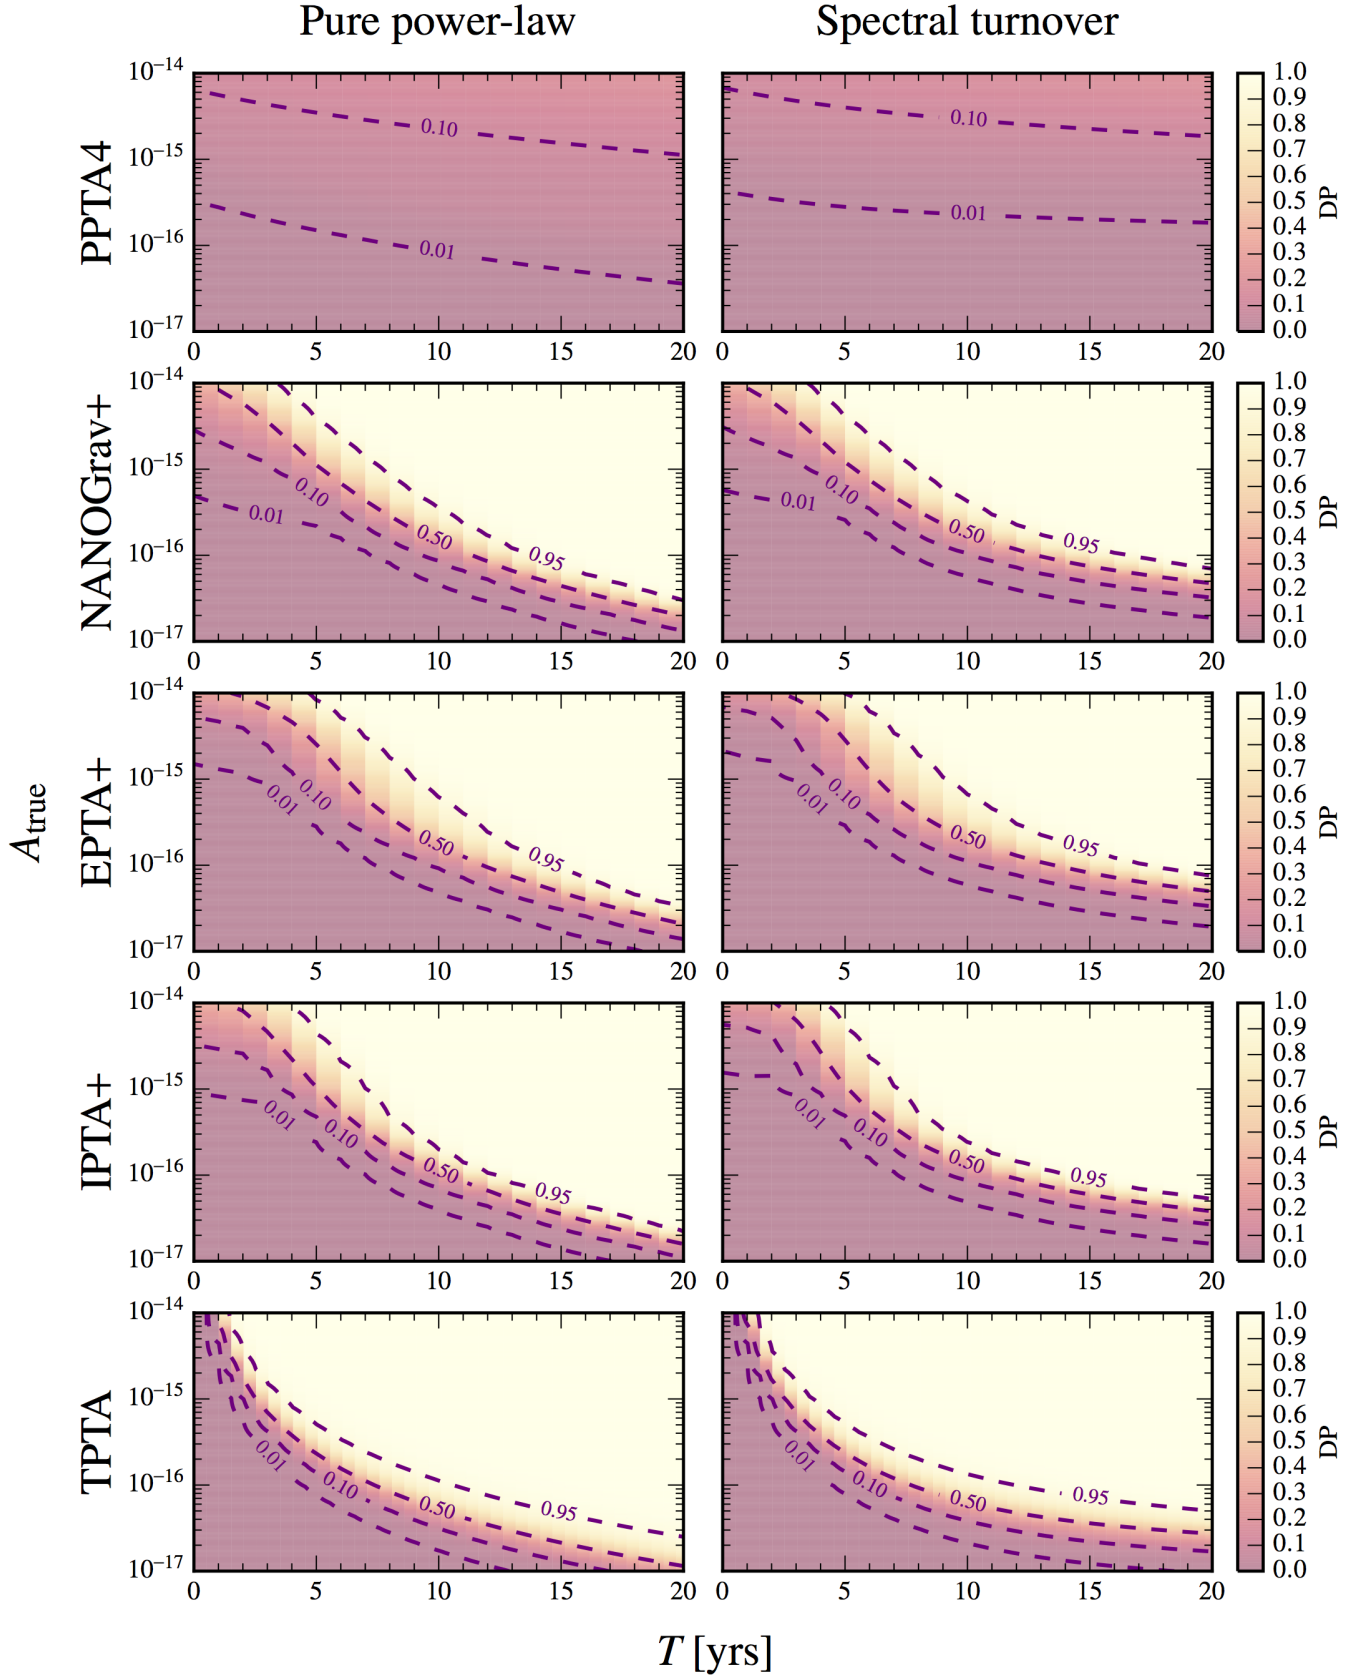


Figure 2. Detection probabilities (DPs) for all arrays considered in this work, as a function of the true GW background amplitude A_{true} and of observation time T beyond the existing dataset. The left panels were derived for pure power-law GW background spectra, and the right panels for turnover spectra with a knee at $f = 1/(11 \text{ yrs})$ due to stellar hardening.




Figure 3. Summary of results for the growth of GWB detection probability (DP) with further observation time in each array. Blue and red lines are for zero and 90% binary stalling, respectively. Dashed lines correspond to the case where the true background spectrum has a turnover at $f = 1/(11 \text{ yrs})$ due to binary stellar hardening.

els to these TPTA pulsars—the DP will already be close to unity when low-frequency timing noise begins to have a significant influence.

We stress that there are several important caveats to our analysis. We have focused on detecting the stochastic background rather than deterministic signals. Even for these, using data from more pulsars is desirable, in order to build evidence for a coherent signal. We have assumed for each pulsar a single value of measurement noise that does not improve over time (as occurs in reality when receivers and backends are upgraded), nor do we consider interstellar medium effects such as dispersion (see, e.g., [Stinebring 2013](#)). The influence of a GWB spectral turnover depends on its frequency, which is a function of the typical environmental properties of galactic nuclei (either directly or in how these properties drive SMBHB orbital eccentricity)—our choice corresponds to reasonable assumptions about the stellar mass density of SMBHB environments. The timelines for the growth of detection in each PTA are approximate, intended to emphasize the differences between PTAs suited to upper-limits versus detection, and to demonstrate the influence of various binary stalling and environmental scenarios on detection probabilities.

We conclude by emphasizing the different demands of placing stringent upper limits on the stochastic back-

ground versus actually detecting it. To wit:

- Highly constraining, astrophysically significant upper limits are achievable with only a few exquisitely timed pulsars, but such a PTA is suboptimal for the detection of a stochastic GW background.
- Timing many pulsars allows for the quadrupolar spatial correlation signature of the SMBHB background to be sampled at many different angular separations, enhancing prospects for detection.
- Adding more pulsars regularly to PTAs will continually improve detection probability, in addition to the gains already made by timing existing pulsars for longer, and will help to mitigate the deleterious influences of binary stalling and environmental couplings.

4. ACKNOWLEDGEMENTS

It is our pleasure to thank Pablo Rosado, Alberto Sesana, Jonathan Gair, Lindley Lentati, Sarah Burke-Spolaor, Xavier Siemens, Maura McLaughlin, Joseph Romano, and Michael Kramer for very useful suggestions. We also thank the full NANOGrav collaboration for their comments and remarks. SRT was supported by an appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory, administered by Oak Ridge Associated Universities through a contract with NASA. MV acknowledges support from the JPL RTD program. JAE and RvH acknowledge support by NASA through Einstein Fellowship grants PF4-150120 and PF3-140116, respectively. CMFM was supported by a Marie Curie International Outgoing Fellowship within the European Union Seventh Framework Programme. This work was supported in part by National Science Foundation Physics Frontier Center award no. 1430284, and by grant PHYS-1066293 and the hospitality of the Aspen Center for Physics. This research was performed at the Jet Propulsion Laboratory, under contract with the National Aeronautics and Space Administration.

REFERENCES

- Arzoumanian, Z., et al. 2015a, ArXiv e-prints, 1508.03024
 —. 2015b, *Astrophys. J.*, 813, 65
 Bassa, C. G., et al. 2016, *MNRAS*, 456, 2196
 Burke-Spolaor, S., & Simon, J. 2015, private communication
 Caballero, R. N., et al. 2015, ArXiv e-prints
 Chamberlin, S. J., Creighton, J. D. E., Siemens, X., Demorest, P., Ellis, J., Price, L. R., & Romano, J. D. 2015, *Phys. Rev. D*, 91, 044048
 Gair, J., Romano, J. D., Taylor, S., & Mingarelli, C. M. F. 2014, *Phys. Rev. D*, 90, 082001
 Hellings, R. W., & Downs, G. S. 1983, *Astrophys. J. Lett*, 265, L39
 Hobbs, G. 2013, *Classical and Quantum Gravity*, 30, 224007
 Hobbs, G., et al. 2010, *Classical and Quantum Gravity*, 27, 084013
 Jaffe, A. H., & Backer, D. C. 2003, *Astrophys. J.*, 583, 616
 Janssen, G., et al. 2015, *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 37
 Kramer, M., & Champion, D. J. 2013, *Classical and Quantum Gravity*, 30, 224009
 Lentati, L., Alexander, P., Hobson, M. P., Taylor, S., Gair, J., Balan, S. T., & van Haasteren, R. 2013, *Phys. Rev. D*, 87, 104021
 Lentati, L., et al. 2015, *MNRAS*, 453, 2576
 Lentati, L., et al. 2016, *MNRAS*
 McLaughlin, M. A. 2013, *Classical and Quantum Gravity*, 30, 224008
 McWilliams, S. T., Ostriker, J. P., & Pretorius, F. 2014, *Astrophys. J.*, 789, 156

- Mingarelli, C. M. F., Sidery, T., Mandel, I., & Vecchio, A. 2013, *Physical Review D*, 88, 062005
- Rajagopal, M., & Romani, R. W. 1995, *Astrophys. J.*, 446, 543
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., Hobbs, G., & Manchester, R. N. 2014, *MNRAS*, 442, 56
- Reardon, D. J., et al. 2016, *MNRAS*, 455, 1751
- Rosado, P. A., Sesana, A., & Gair, J. 2015, *MNRAS*, 451, 2417
- Sampson, L., Cornish, N. J., & McWilliams, S. T. 2015, *Phys. Rev. D*, 91, 084055
- Sesana, A. 2013, *MNRAS*, 433, L1
- Sesana, A., Haardt, F., Madau, P., & Volonteri, M. 2004, *Astrophys. J.*, 611, 623
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, 390, 192
- Shannon, R. M., et al. 2015, *Science*, 349, 1522
- Siemens, X., Ellis, J., Jenet, F., & Romano, J. D. 2013, *Classical and Quantum Gravity*, 30, 224015
- Stinebring, D. 2013, *Classical and Quantum Gravity*, 30, 224006
- Vallisneri, M. 2012, *Phys. Rev. D*, 86, 082001
- van Haasteren, R., & Levin, Y. 2013, *Monthly Notices of the Royal Astronomical Society*, 428, 1147
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2009, *MNRAS*, 395, 1005
- Verbiest, J. P. W., et al. 2016, *MNRAS*
- Wyithe, J. S. B., & Loeb, A. 2003, *Astrophys. J.*, 590, 691